

Much of the communication between applications or from an application to an end user is in the form of text files. Examples are the traditional batch interface files that many "legacy" applications still use, and more contemporary formats such as html, xml and rtf. The latter are examples that indicate that text file interfaces will be with us for a long time.

The problem is that dealing with these files requires a great deal of application logic to parse the incoming text file and do something with the content, or to assemble application data and format it into an appropriate configuration.

There are three basic problems in text file processing. The first is how to describe and program for the structure of the text file. The second is how to achieve a mapping of data between the text file and the application (note that it assumed here that the application has some independent means, in this case EADP, of moving the data from some stored format such as a database). Inherent in this is the third problem: describing the flow of control needed to process the file. The structure of the text file inherently describes the sequence of processing needed to process it; however, organizing this structure into application logic is often a major source of application programming.

Summary Of The Invention

An object of the present invention is to improve text file processing.

Another object of this invention is to provide an improved text file interface support in an object oriented application.

5 A further object of the present invention is to handle the problems of pattern matching and flow of control, in text file processing, by creating a structure of templates and macros that correspond to the structure of the text file to be processed.

10 These and other objects are attained with a method and system for processing a text file in a computer application. The method comprises the steps of forming a template from fragments of the text file; using
15 the template as an overlay for parsing incoming files, or as a prototype to generate a segment of an output file; and using a macro class to map data from the text file to an application. The macro class is embedded as a keyword within the template, so that when the template reaches
20 the keyword, the template calls the macro class to further process the text file.

25 The macro class may be used to read in a segment of the text file and to use the segment to initiate application update processing. Also, the macro class may be used to derive data from the application and to format it into the text file. Preferably, an interface controller is provided to prevent structure clashes by placing text
30 data into appropriate places in a complex object structure as the text file is processed.

Figure 2 shows major features of a preferred embodiment of the invention.

Figure 3 is a flow chart illustrating an aspect of an embodiment of the invention.

Detailed Description Of The Preferred Embodiments

As mentioned above, there are three basic problems in text file processing. The first is how to describe and program for the structure of the text file. The second is how to achieve a mapping of data between the text file and the application (note that it assumed here that the application has some independent means, for example EADP, of moving the data from some stored format such as a database). Inherent in this is the third problem: describing the flow of control needed to process the file. The structure of the text file inherently describes the sequence of processing needed to process it; however, organizing this structure into application logic is often a major source of application programming.

As generally outlined in Figure 1, the present invention effectively addresses each of these problems. More specifically, this invention solves the first problem, describing the text format, by using fragments of text (templates) as overlays for parsing incoming files, or as prototypes to generate segments of output files. There are several important advantages to this approach. The templates are isolated from all other application logic, so it is easy to see why a particular template is producing a resulting text file. The template is a

literal image of the text fragment it processes, so it is possible to create the templates from samples of the text file. If the format changes, it is easy to change the corresponding template. For incoming files, there is an important advantage to being able to use a template as a mask to parse the text file. This replaces hard to decipher application logic typically used for parsing.

The second problem of mapping data from the text file to the application is solved by specialized macro classes. These come in two flavors (one for input and one for output). The input macro read in a segment of the text file and use it to initiate application update processing. The output macros derive data from the application and format it into the text file (special format classes are used to describe how to transform the output into the proper shape). This invention provides a basic set of macros, and the facilities to add more as needed.

The solution to the third problem, flow of control, is solved by the interaction of templates and macros. A macro is embedded as a special keyword within a template. When the template reaches that point, it calls the macro. The macro in turn is passed another template name as part of its invocation; as a part of its processing it can invoke that template (which in turn calls other macros, etc.). This nested aggregation of templates and macros allows a processing structure to be built up that mirrors the inherit structure of the text file. Since the behavior of the macro depends both on its internal logic and the template it is passed to invoke, it is possible

to reuse the same macro to do different things by passing it a different template. The net effect is that the bulk of the logic needed to describe flow of control is included in the template structure. The template/macro combination is the ultimate expression of the ideal of letting the target file structure determine the structure of the application needed to process it.

The above procedures may produce "structure clashes"; this is a fourth problem that needs to be solved within this context, particularly for incoming text files which must be mapped to application data.

This invention includes an interface controller which sets up a temporary complex object structure in memory. The interface macros pigeonhole data into appropriate places in the structure as the text file is processed. After the entire file is processed, the internal structure is used to process updating data into the application. This master processing module can be used for any EADP compliant application, as can the base interface macros that are provided as part of this invention. The great advantage here is that processing a new batch interface for a new application becomes mostly a matter of setting up the template structure to map out the text file. The macro processing and the base application processing are already in place.

With reference to Figure 2, the preferred embodiment of this invention includes the following major features.

1. Templates that support variable substitution and

conditional or iterative generation for output files;

2. A mechanism to fill in variable substitutions in the templates;

5

3. A mechanism to handle iterations and conditional processing;

10

4. A flexible template processing mechanism which operates within an object oriented environment;

5. A flexible and easily extended formatting mechanism for output;

15

6. A mechanism for template creation and editing;

7. Special support for complex object presentation;

20

8. Templates that support parsing of the input file and mapping of its contents into a complex object structure suitable for update processing;

9. A mechanism to extract data from the input stream;

25

10. A mechanism to map the structure of the input file into a complex object structure suitable for internal processing;

30

11. A mechanism to organize the updating data from the input file and to apply it within the application;

12. A flexible and easily extended formatting mechanism;

and

13. A mechanism for input template creation and editing.

- 5 1. Templates that support variable substitution and conditional or iterative generation for output files.

10 The templates are fragments of text that include the literals for the output text stream. These capture the predictable pattern of the output stream. Each template may include substitution points, which are to be filled in with application data, or calls to macros to perform conditional or iterative processing.

- 15 2. A mechanism to fill in variable substitutions in the templates.

20 A simple variable substitution in the template is indicated by a tag name enclosed by dollar signs. When the template is processed, a Java Hashtable is prepared with the tag names and substitution variables. These are matched against the tag names in the template as the output text stream is created.

- 25 3. A mechanism to handle iterations and conditional processing.

30 The name of a Java class can be used as a substitution variable. If a class name is used, the "macro" method on that class is invoked (it is passed other information which can be included in the substitution variable, such as the name of the next template to process).

Iterations, conditional logic, and the preparation of data for the next template is handled by these "macro" classes.

- 5 4. A flexible template processing mechanism which operates within an object oriented environment.

There are two aspects to this.

- 10 a. "Macro" classes which all inherit from a common class that defines the "macro" method. There are two base macro classes (one for input and one for output). An interface controller keeps track of the macros needed for a particular process so that they do not need to be
15 continually reinvoked using Class.forName.

- 20 b. Java classes which controls template processing. The class for output processing handles functions such as parsing the templates to find substitution variables, resolving the substitutions, and invoking macro classes as indicated in the substitution variables. The class
25 for input processing handles functions such as parsing the templates to find keywords, synchronizing the templates with the input file, preparing a memory structure with the updates, and then applying the updates to the database.

- 30 5. A flexible and easily extended formatting mechanism for output.

Much of the format of the output stream is determined by the literal text within the templates themselves. In

addition, the name of a formatting class can be included in the variable substitution name. These formatting classes all inherit from a common class which defines the "format" method. This method can be redefined to do any specialized formatting required for that substitution variable in that template. Of course, the same formatting class can be reused in other places.

6. A mechanism for template creation and editing.

The templates for each application are stored as a Hashtable that is associated to that application's definition class (as described in copending application Complex Object Navigation, Presentation and Manipulation). The template editor is implemented as a Java bean custom editor.

7. Special support for complex object presentation.

A specialized macro is provided which supports the navigation of a complex object structure. This means that only the templates need to be provided to give different types of reports or batch interfaces based on the data in the database underlying the complex object structure.

8. Templates that support parsing of the input file and mapping of its contents into a complex object structure suitable for update processing.

The templates are fragments of text that include the literals for the input text stream. These capture the

predictable pattern of the input stream. Each template includes macros which specify points in the input stream from which data is to be taken (and what to do with that data). The macros can also specify that the field is of a fixed length, so that that number of bytes is automatically advanced in the stream.

9. A mechanism to extract data from the input stream.

An extraction point in the template is indicated by a keyword (a set of characters enclosed by dollar signs). Parameters within the keyword are separated by commas. The first parameter is the name of a macro class. However, the macros used for interface process implement a macro method with slightly different parameters, as described below. The additional parameters give information such as the name of the field to be extracted, which database table it belongs to in the current application, and possibly the length of the input field. Two strategies to recognize the boundaries of the input fields are supported:

a. Some input files use separators between fields and the fields have variable lengths. In this case the separator would be a literal in the template.

b. Others use fixed length fields with no separators. In this case the length of the field is specified in the keyword.

10. A mechanism to map the structure of the input file into a complex object structure suitable for internal processing.

5 The complex object structure of the target application is defined using the customization mechanisms described in copending application Complex Object Navigation, Presentation and Manipulation, Each keyword in the template used for data extraction specifies the column
10 name and the internal class name for the data to be extracted (specialized keywords can be constructed to assign the same data to more than one table or column). This allows a rapid and simple way to map the contents of the input file onto the internal structure of the
15 application. It also gives considerable flexibility, since the contents of one interface record may be reorganized to map to several internal classes (and thus different tables in the target database) or vice versa.

20 11. A mechanism to organize the updating data from the input file and to apply it within the application.

25 This invention takes advantage of the customization facilities described in copending application Complex Object Navigation, Presentation and Manipulation. That invention shows how each database table is associated to an "internal class" that controls updates to that table. The parsing information within the template specifies for each field which internal class it should be assigned to,
30 and the database column name for the field. A memory structure is created to organize this data by internal class and by key data. Within this structure, the

internal classes are organized by complex object structure. When the updates are applied to the database, this is done using the same methods that apply data interactively. The complex object structure is followed, so that rulers are updated before subobjects. All the data for subobjects is linked in memory to the data for the ruler, so that if the ruler update fails, update of its associated subobjects can be avoided.

12. A flexible and easily extended formatting mechanism.

This is similar to the formatting strategy for output files, but in reverse. The format class is used to translate the string data from the input field into a suitable format for internal processing. This can include such things as skipping quotes or extra blank in the input file.

13. A mechanism for input template creation and editing.

The Java implementation adds an interfaceDictionary property to the EADPTemplateDefinition class as a place to define the interfaceTemplates for a particular application. This is then customized for the database definition class for a particular database. The templates for that database are used when an interface for that database is processed.

Each record in the interface file must be matched to a unique template. The type of input record for the template can be specified in two ways:

a. The template name can start with a table name. This table name is then specified when the interface file is processed. Only templates that begin with this name will be used.

5

b. If the record has an identification symbol at the beginning of the record, this can be used to qualify the name of the template used to parse that record. This is useful when the interface file has a mix of record types. An asterisk in the template name indicates that it can be used for any record in the interface file.

10

The discussion below describes in greater detail several of the above-mentioned features of the preferred embodiment of the invention.

15

Output text generation

1. Template definition

20

The syntax of templates is kept very simple so that templates can be easily created from samples of the target output format. A template is a piece of text with imbedded keywords. The keywords are delimited by dollar signs. For example:

25

text1.. \$keyword1\$ text2.... \$keyword2\$... etc.

A template can span multiple lines of text. However, a special keyword \$+\$ at the end of a line indicates that a new line character should not be added when the template is processed.

30

When a template is processed, the text outside the keyword is passed along as is into the output stream. The text inside a keyword is evaluated according to the following rules:

5

a. The text up to the first comma is checked to see if it matches the name of a VisualAge class. If it does, the "macro" method of that class is invoked. The remainder of the text within the keyword (after the comma, up to the dollar sign) is passed as a parameter to the macro method. It is then passed to the receiving method (additional commas may be used to separate information).

10

For example, the receiving macro can be passed the name of the next template to process using this technique. This is useful when the output stream requires a nested sequence of templates and macros (for example, a report of purchase orders and line items for each purchase order).

15

20

b. If the first part of the text is not the name of a class, it should match the name of a variable in the substitution list which was prepared for use by the template when it was invoked. In this case, any text after the comma is assumed to be formatting information:

25

i. The first piece is the name of the class which will do the formatting (for example PadRightFormat). If this is omitted, no special formatting is done. The data appears just as it would on the list or entry panels.

30

ii. If there is more data (delimited by a second comma)

this is passed to the formatting class. For example, the length of the field is passed to PadRightFormat.

2. Resolution of variable substitutions in the templates.

When each template is processed, it is presented a Java Hashtable which has as keys the names of the variables (matching the variables names specified within the template). The values are substituted into the template (they may be modified by the specified formatting routine). The variable substitutions are prepared by macros prior to template invocation.

3. A mechanism to handle iterations and conditional processing.

This invention avoids creating another programming language by adding logic constructs within the template itself. The only branch to logical processing is through invocation of macros, which are coded in Java (and can take advantage of the full power of that environment).

Since templates can invoke macros and macros can in turn invoke templates a recursive chain of macros and templates can be used to handle nested structures within the output stream (for example, reporting orders, and within each order, all the line items for that order).

4. Macro classes

Macro classes all inherit from a common class, com.ibm.eadp.macros.EADPMacroBase, and redefine the

instance method macro. The parameters passed are as follows:

a. segment

The macro was invoked from with a template because the keyword (text within dollar signs) contained the name of the macro. There may be other text after the macro name, and this is passed as the segment parameter. The various pieces of information in the keyword are separated by commas. Information that is commonly included here is the name of the template the macro should use for its processing. This allows reuse of the same macro to process different templates.

b. currentClass

This is a reference to an instance of a Java class. Messages can be sent to that class to derive data to populate the variable list.

c. varList

This is the variable list (Hashtable) that was presented to the calling template. Data from it may be used by the macro to build its own variable list.

d. outStream

This is the output stream for the text stream that is to be generated (of type StringWriter).

In addition, the macro class has as a property the current controller, which is the instance of EADPTemplateController that is processing the templates (and which invoked that macro).

5

5. A Java class which controls template processing.

Template processing is controlled by com.ibm.eadp.macros.EADPTemplateController. It is expected that when a new instance of controller is created, it will be assigned a data base definition class (a child of EADPDatabaseDefinition as defined in Complex Object Navigation, Presentation and Manipulation). EADPDatabaseDefinition has an additional property, templateDefinition, of type com.ibm.eadp.macros.EADPTemplateDefinition. This in turn has two properties, interfaceDictionary and templateDictionary, both of type com.ibm.eadp.macros.EADPTemplateDictionary. This class is a child of Hashtable (redefined to give it a custom editor).

When a new instance of EADPTemplateController is created to be used for processing, it is assigned a database definition class (this gives it access to the database) and a template dictionary (typically the templateDictionary property of the templateDefinition for the database definition).

Some methods:

a. getTemplateDictionary

This method has as its first parameter the key for the variable substitution. This is used to find the value in the varList, which is a dictionary containing the variable substitutions to be used when processing the template.

The second parameter contains formatting information (separated by commas). The first segment is the name of the formatting class. The remainder is passed as the "pattern" to the format method on the formatting class. If there is no formatting information specified (the format parameter is a null string), the EADPBasicFormat class is used. Format classes are resolved using the same strategy as the macro classes (and formatForName methods).

6. A flexible and easily extended formatting mechanism.

All formatting classes inherit from com.ibm.eadp.macros.EADPBasicFormat class and redefine the format(string,pattern, outputStream) method. The first parameter here is the string to be formatted. The second parameter may contain additional formatting instructions, separated by commas (each child would add its own logic to parse the formatting instructions). The formatted result is appended to the passed outputStream.

The base method just appends the passed string to the outputStream without any additional formatting.

7. A mechanism for template creation and editing.

The technique to create the Java initialization string and to initialize the string is similar to the one described in the above-identified copending application "Flexible Help Support In An Object Oriented Application." Only two separators are needed (for the key and text of each template). New line markers are added to the generated string at each line of the template so that the generated code is readable; an extra new line marker is also added within the initialization string for each template line so that the initialized text has the proper line breaks.

8. Special support for complex object presentation.

When a macro is processed, the remaining information in the keyword is passed as the first parameter. The macros described here expected the keyword to contain the name of the internal class (as defined by the above-identified copending application "Complex Data Navigation, Manipulation And Presentation Support For VisualAge Java") and the name of the next template to process. The "current class" would be an instance of the ruler class (for example, the ruler class would be for orders and the name of the class in the keyword would indicate line items for that order). The macro classes have a currentController property which is assigned by macroForName as the macro instance is created. This instance is the one that gets the processTemplate call.

The ReportMacro class redefines the macro method. As mentioned above, the first parameter should contain a string which has the next internal class name and the

next template name, separated by commas. The current class should be an instance of a child of EADPApplicationClass. For the first level, the class name of a top object can be passed, along with an the database definition class for the process. If this is done, the managerForName method on EADPDatabaseDefinition is used to find an instance of the class. Otherwise, an instance of a child of EADPApplicationClass (the ruler class) is passed as the current class, and its subManager method is used to find the subobject class (using the current row of the ruler class as the ruler row).

The class name and template name are parsed out of the first parameter. If the class name matches the name of the current class (this would be true for the first level of the process), the current row of the current class (as the first parameter) and the current class (as the second parameter) are passed to the processRow(row, class, template, outStream, varList) method. Otherwise, a new instance of the type of class indicated by the class name is created (this should be subobject of the current class) and is assigned the current class as a ruler. The current row of the current class is used to open the subobject class (this finds all the subobject rows that match the current row of the ruler, e.g. all lines for a particular order). For each row of the subobject class, the processRow method is called (passing the subobject class as the second parameter). In both cases, the template name parsed from the keyword, and the passed varList and outStream are passed as the additional parameters.

The processRow method copies the passed varList into a new dictionary, and the uses the row dictionary for the passed row to add additional entries. The values in the row dictionary are converted to string format using the getStringValue method on EADPPersistentObject, which does the property editor lookup to format the string value. The processTemplate method on TemplateMetaclass is then called, passing the template name, the class (which is now the subobject class), and the new variable list. Combined with a set of templates that call out various pieces of the complex object structure and provide formatting information, this allows navigation down through the levels of the complex object.

Input text processing

All methods are instance methods. To avoid creating many instances of the same macro class, the classes are cached and reused. The setStringValue method on EADPPersistentObject is used to convert the data from string format to internal format, and this is done as the data is applied to the database. Before that, all data is held in string format. The setStringMethod uses the editor dictionary for the row's data manager to find the right property editor to convert the value to object format.

One difference with the data update function described in the above-identified copending application "Complex Data Navigation, Manipulation And Presentation Support For VisualAge Java," is the way a prototype for a new row is created. In the function described in the above-

identified copending application "Complex Data Navigation, Manipulation And Presentation Support For VisualAge Java", a new row is being added to rows for a result set that has already been created, so that the metadata for the result set can be used to determine the data type for the columns. Here, a new row has to be defined before a query has been issued, so database metadata has to be used. Since this is an expensive call, and a typical interface will not involve many different types of rows, the results are saved and reused.

Several classes and methods are described below:

1. EADPInterfaceView

This defines the visual part used to process a batch interface file. It is customized to provide an interface view for a particular application by customizing its database definition property. This class inherits from Panel, and it designed to be included as a visual bean in the interface view for a particular application.

It has the following features.

a. Table name text area

This is used to specify the table name prefix.

b. Text area for the interface file

The interface file is presented in a text area which

allows editing. This can be useful when the file that is sent to be processed has header and trailer records which are not needed as it is being processed. They can be stripped off manually before the interface is processed.

5

c. File handling beans and buttons.

The EADPTextHelper class defines the readTextFromFile and writeTextToFile methods. These methods are connected to the Open File and Save File buttons. Standard file dialog beans (java.awt.FileDialog) are used find directory and file names.

10

d. Continuation

15

This text area is used to specify a continuation character which may be used by the interface templates.

e. processInterface

20

This method is invoked when the Process Interface button is pressed. The parameters passed are the interface file, the table name, and the continuation character. To begin processing, a new instance of

25

EAPDInterfaceController is created, and it is assigned the database definition property as its database definition. The interfaceDictionary property of the templateDefinition property of the database definition is assigned to the templateDictionary property of the controller (this is how the controller knows which templates to look for). The datastore property of the database definition is used to connect to the database.

30

Next, the processInterface method on EADPInterfaceController is called to parse the interface and apply the changes.

5 f. currentDefinition property

This property is set to the database definition for the application when the visual bean is customized. This is how the Java version knows which database to process.
10 The database definition class also has the connection information that is needed to establish a connection to the database to begin processing.

2. EADPInterfaceController

15

This class has methods that process the interface file.

a. processInterface

20

This is passed the interface file (as a String), the table name, and the continuation character.

25

This method controls the interface processing. It creates a Hashtable (the processDictionary) which acts as a storage area in memory for the updating data. This dictionary is passed as a parameter to all the methods (including macros) that are used to process the interface file.

30

The first step is to break the interface file into records. The continuation character is used to combine lines from the input file into logical records. For each

record, the matching template is determined by calling the templateForSegment method. The record and its associated template are then passed to processTemplate method.

5

Once the entire interface has been processed, the updates are applied by calling the doApply method. This iterates through the process dictionary and calls apply.

10

b. templateForSegment

This method is passed the current line of the interface file and the table name specified as the interface was initiated.

15

It finds the template that will be used to parse the current record. It uses the templateDictionary property (assigned as the controller was initiated) to locate the dictionary of templates for the current application. It then iterates through the keys of the dictionary, looking for ones that start with the passed table name. If a match is found, it then checks the remainder of the key. If it is an asterisk, the template is used without checking it against the input record (an asterisk indicates that the template is to be used for all records for that table type). Otherwise, the remainder of the template name is checked to see if it matches the beginning of the record (this is used if the records start with flags that indicate the record type). The template is then returned to the caller.

20

25

30

c. processTemplate

The parameters passed in are the template that was found as described above, the current record of the interface, and the processDictionary.

5 This method first creates a scratch pad entry in the dictionary for holding the processing results for the current row. This is a new dictionary added at the key value 'currentRow' (the processDictionary is a nested set of dictionaries, so most of its entries are other
10 Hashtables). This will be referred to as the currentRow dictionary.

15 Next, the template is broken up into individual lines (for convenience, the template may consist of multiple lines. This allows each field of the record to be described by a line of the template, which makes it easier to describe and understand the record structure. Typically any existing documentation of the record structure will have this format, and this makes it easier
20 to convert that documentation into a template to process the record.

25 For each line of the template, the processSegment method is called. This will create entries in the currentRow dictionary.

30 Once all the lines have been processed, the updateDictionary method is called (passing the processDictionary). This will move entries from the "currentRow" dictionary to their final position in the memory structure.

d. processSegment

The parameters passed in are the current line of the template that was found as described above, the current record of the interface, and the processDictionary.

This method finds a keyword in the passed line of the template by looking for text enclosed by dollar signs. The fragment of text between this and the next keyword is then found (this is the trailer). These are passed to processKeyword.

Next, the method looks for any characters outside dollar signs, and advances the position in the interface record to match those characters (the upTo method in EADPStringReader is optimized to do this). This is how separator characters in the interface record are described in the template. The separator is included as literal text after the extraction keyword.

This process continues until the end of the template line is reached.

e. processKeyword

The passed parameters are the keyword, the fragment of template after that keyword, the interface record, and the processDictionary. The macro name is parsed from the keyword (it is separated from the rest of the keyword by a comma), and is used to find the macro class. The macro method is called on the macro class. Typically, the EADPColumnInterfaceMacro would be called at this

point. However, in some cases the same column data needs to be used in several places (for example, a row in the input may correspond to both a ruler and subobject in the target application if the source application was not well structured).

f. macro(keyword, trailer, segment, dictionary) method in EADPColumnInterfaceMacro

This method will be described now to make it easier to understand the flow of control. Keep in mind that other similar macros could have been invoked at this point (for example, the EADPTwoColumnInterface which places the same column data into two different internal classes). The macro invocation is determined by the placement of the macro name in the template, so the template structure is driving the flow of control at this point.

The parameters passed are the remainder of the keyword (which contains parameter information for the macro), the fragment of the template line between this keyword and the next one, the current interface record, and the processDictionary.

The following parameters are contained in the remainder of the keyword, separated by commas:

1. Internal class name

The name of the internal class to process the update. This may be omitted, if a keyword specifying a standard internal class for the template has been provided (this

the format method of the formatting class.

Once the data has been extracted and formatted, it is placed into the memory structure defined by the passed processDictionary. The dictionary entry at the key "currentRow" is accessed. This as a dictionary keyed by the name of the internal class used to hold the data. The entry for the internal class passed in the keyword is found (if none is present, one is created). This entry is a dictionary that holds data for the row (the keys are the column names). The column name passed as a parameter in the keyword is used to add the data just extracted from the interface record. The values are stored in string format (they are converted later during processing in the apply method).

g. updateDictionary

This method moves rows from the "currentRow" dictionary to their permanent position. To find the permanent position, the key for the row needs to be known (this is why this process is delayed until the entire template has been processed against the interface record. The key columns may be positioned anywhere within the interface record, so the entire record needs to be parsed before it can be classified).

The method iterates over the entries in the currentRow dictionary. The keys are internal class names, and the values are row descriptions (dictionaries of column names and values). These entries are passed to the hasKeys method to determine if the key columns were found. If

so, the addRowToMaster method is called to place the row data in the proper place in the memory structure.

h. hasKeys

5

This method is passed a key and value from the currentRows dictionary. The key is the name of an internal class. That is used to find the complex object node (EADPComplexObjectNode) for that class in the classDirectory property of the complexObjectStructure property of the currentDirectory property of the database definition (this is encapsulated in the nodeForName method in EADPDatabaseDefinition). The keyCols property of the complex object node is then used to make sure that each has a corresponding entry in the passed value (which is a dictionary of column values keyed by column name).

10

15

i. addRowToMaster

20

This method is passed a key and value from the currentRows dictionary. The key is the name of an internal class. The passed value is a dictionary of column values keyed by column name.

25

The master dictionary structure is a set of nested dictionaries. It follows the complex object structure, with each instance of a ruler object controlling lower level dictionaries of its subobjects. To achieve this, a two level nesting of keys is used. This first level key is the name of the internal class. The second level is the key information for a particular instance of that class (this is stored as a Vector of the string

30

values for the key in the same order as the keyCols attribute that was used to set it up). For example, if the complex object structure consists of orders and line items, the master dictionary structure might look as shown in Figure 3.

To accommodate this structure, a node class EADPInterfaceNode is used. It consists of two attributes, rowDictionary the dictionary of column values for the row, and subobjectDictionary, a dictionary of EADPInterfaceNodes keyed by the subobject class names.

The dictForRow method is used to find the proper position within this structure. This positions to the entries for the internal class that are underneath the rulers specified by the key information for the row. Next, a check is made to see if there is an entry for key data matching the current row. If not, a new entry is created. The row data at the entry is then updated with data from the passed row value dictionary.

j. dictForRow

The parameters are the row dictionary, the internal class for the row, and the processDictionary.

This method calls the rulersForClass method to build a Vector of EADPComplexObject nodes for the rulers for the internal class. It then iterates through the Vector. To begin the iteration, the current dictionary is set to the processDictionary, and the current class is set to the first entry in the ordered collection (the top level

ruler).

1. The current class name is used as a key in the current dictionary. The corresponding value is a "row key
5 dictionary" of instances of EADPInterfaceNodes, keyed by the key values for instances of the current class (for example, if the top level is Orders, these would be keyed by the order number).

10 If there is no entry at the key, a new dictionary is created.

2. Once the row key dictionary has been found, the key
15 for the current class (derived from the values in the passed parameter for row data) is used to find an entry that matches the key values for the row being stored. The entry will be an instance of EADPInterfaceNode. If none is found, a new entry is created.

20 3. The current dictionary is set to the subobjectDictionary attribute of the EADPInterfaceNode.

4 .The current class is advanced to the next entry in the
25 Vector.

When this iteration completes, the current dictionary
will be set to the subobjectDictionary attribute of the
EADPInterfaceNode for the immediate ruler of the passed
parameters (it will remain at the main dictionary if the
30 passed internal class had no rulers).

k. doApply

This method is called after the interface has been read into the internal memory structure. It is passed the processDictionary. It makes the initial calls to apply. The entries in the processDictionary are passed (the key is passed as the internal class name and the value is the dictionary of interface nodes). At this level, the passed ruler class and ruler node are null.

1. apply

The passed parameters are the name of an internal class, a dictionary of EADPInterfaceNode instances, the ruler class, and the ruler row.

The method invokes itself recursively to work down through the nested structure of dictionaries and interface nodes. At each level it iterates through the dictionary values. Each entry is an EADPInterfaceNode. The rowDictionary attribute holds updating information. An instance of the data manager for the internal class is created using the internal class name. If the passed ruler class is null, the managerForName method on the definition class is used. If the ruler class is not null, the subManager method on its data manager is used, passing the ruler row as the second parameter (this will automatically set up the ruler list). Both these techniques pass along the connection information so that the newly created data manager has database access.

The passed dictionary is a row value dictionary, where the keys are key values (the vector of key values converted to string format). For each entry the key

CONFIDENTIAL

information in the key for the dictionary entry is converted back to a vector of strings (using the `stringToVector` method) and the is passed to the `rowForKeys` method on the data manager (`EADPDataManager`), which is used to find a matching row on the database (if there is none a new blank row is created). If a new row is created, the `fullCompoundType` method on `EADPDAManager` is used to find the column names and types (using database table metadata). Since this is an expensive call, the results are cached by row type in the interface controller.

The editors dictionary is also set up at this time, based on the type of the column. Next, the `rowDictionary` is used to update column values (using the `setStringValue` method on `EADPPersistentObject`, which uses the `editorsDictionary` property of the `datamanager` to find the right property editor for each column and convert the string value to object value using the property editor). The updates are applied to the database (this is done using the facilities of EADP internal processing so that any verifications that were specified during customization are invoked at this time). If no updating information was provided, and the row was not found, processing stops for this leg because the associated subobjects do not have a ruler. Also, if the apply failed because of errors, processing stops for this leg.

If processing can continue on this leg, the method iterates over the `subobjectsDictionary` attribute of the `EADPInterface` node. For each entry, the apply uses the key for the entry (the internal class name of the

subobject), the value of the entry (a dictionary of EADPInterfaceNodes keyed by key values for the subobject) the application class for the data manager, and the row just processed.

5

3. EADPInterfaceBase

10 This class is the base macro class for interface macros and defines the macro(keyword,trailer, segment, dictionary, method).

4. EADPColumnInterface

15 This class was described during the explanation of processing of EADPInterfaceClass. It is used for the majority of column extraction processing.

5. EADPInterfaceFormat

20 This is the basic formatting class, and it defines the format: pattern: method. At this level, the method just returns the input string unchanged.

6. EADPTableType

25

This macro class is used to avoid placing the same internal class name in each keyword within a template. It stores the passed internal class name in a work area in the processDictionary, so it can be used as a default.

30

As mentioned above, the problem of handling text files, both for input and output, has been around for a long

Also, this invention does not require the introduction of a new procedural language within the tool to handle flow of control. All the macros are written in Java. So another advantage is that all procedural logic is written in a standard way. Also, enough information is passed to the macro classes so that they have full access to all the power of the EADP internal classes. This means that the macros can do quite a bit of processing if that is required.

The present invention has been implemented in the Enterprise Application Development Platform (EADP). The user manual for this facility is included herein as Appendix A.

While it is apparent that the invention herein disclosed is well calculated to fulfill the objects stated above, it will be appreciated that numerous modifications and embodiments may be devised by those skilled in the art, and it is intended that the appended claims cover all such modifications and embodiments as fall within the true spirit and scope of the present invention.